

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau




(43) International Publication Date
8 January 2004 (08.01.2004)

PCT

(10) International Publication Number
WO 2004/004245 A1

- (51) International Patent Classification⁷: **H04L 12/56**
- (21) International Application Number:
PCT/IB2002/002492
- (22) International Filing Date: 28 June 2002 (28.06.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant (for all designated States except US): **NOKIA CORPORATION** [FI/FI]; Keilalahdentie 4, FIN-02150 Espoo (FI).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **KOISTINEN, Tommi** [FI/FI]; Tuurakuja 4 A 8, FIN-01600 Vantaa (FI). **HALMET, Karo** [FI/FI]; Lapinlahdenkatu 23 A 12, FIN-00180 Helsinki (FI). **LEPÄNAHO, Henrik** [FI/FI]; Hiihtomäentie 16 E 48, FIN-00810 Helsinki (FI).
- (74) Agent: **LESON, Thomas, Johannes, Alois**; TBK-Patent, Bavariaring 4-6, 80336 München (DE).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: **LOAD BALANCING DEVICES AND METHOD THEREFOR**



(57) Abstract: A method of balancing the load of resources in a packet switched connection within a communication system, said system comprising processing units (11; 21) for performing communication, at least one load balancing unit (12; 22) for distributing the load to said processing units (11; 21), and a data storage (14; 24), said method comprising the steps of: obtaining a current connection state as well as a current load state of said processing units from said data storage (14; 24); selecting by said load balancing unit (12; 22) a processing unit on a per-packet basis; and maintaining information about the load state of each processing unit (11; 21) so that said selecting step is performed by selecting a processing unit to serve and process a respective packet based on the load state.

Load balancing devices and method therefor

Field of the Invention

5 The present invention relates to an automated load
balancing and substantially streamlined resource
management in a communication system. The present
invention is particularly applicable when high-speed
packet based links are interfaced with signal processing
10 resources.

Background of the Invention

Typically, in present communication systems a resource
15 management functionality plays a significant role in any
network element providing heterogeneous signal processing
services. This means that several processing units must
be dedicated to provide the resource management
functionality. Also, quite a high amount of the internal
20 communication bandwidth must be reserved for exchanging
the resource management related control messages.
However, increasing the number of the media processing
units may cause the resource manager to become a
bottleneck that may mean a reduced overall
25 cost-efficiency, and, in practice, longer latency times
when responding to new service requests.

Fig. 3 depicts an arrangement where a load balancing
device 32 is used in a conventional way. Reference
30 numeral 33 designates a device for routing packets of a
communication connection. Particularly, according to this
prior art, the processing unit is selected out of a
plurality of processing units 31 on a per-connection
basis.

35

- 2 -

However, the prior art described above suffers from the following drawbacks. This conventional resource management system is not scaleable which means that there is a large number of dedicated units as well as a large
5 number of control messages. In addition, delays in responding to service requests occur. Moreover, the utilization of statistical multiplexing inflicts difficulties.

10 Particularly, the processing times in the packet switched connection according to the prior art are not deterministic. Specifically, in the conventional arrangement, one digital signal processor (DSP) is receiving and processing several channels (e.g. 8...16)
15 simultaneously. This means that a packet must wait a non-deterministic time before it is processed, that is, depending on how many channels are processed before that specific channel.

20 Summary of the Invention

Accordingly, it is an object of the present invention to overcome these shortcomings of the prior art and to provide a streamlined and cost-effective way to manage
25 the available resource pool.

According to the present invention, the object is solved by providing a method of balancing the load of resources in a packet switched connection within a communication
30 system, said system comprising processing units for performing communication, at least one load balancing unit for distributing the load to said processing units, and a data storage, said method comprising the steps of: obtaining a current connection state as well as a current
35 load state of said processing units from said data

- 3 -

storage; selecting by said load balancing unit a processing unit on a per-packet basis; and maintaining information about the load state of each processing unit so that said selecting step is performed by selecting a processing unit to serve and process a respective packet based on the load state.

Here, the data storage can be accessed to by said load balancing unit or said processing units. Further, the information about the load state may be maintained as a Boolean state, i.e. to indicate free or not free.

The selection of a processing unit can be done in a round-robin fashion.

Further, a supported service profile for each processing unit can be maintained in addition. In this case, the supported service profile can be used as additional selection criteria.

In the method according to the present invention, the load balancing unit can obtain a load state from each processing unit upon a hardware based mechanism or a packet based mechanism. In the latter case, a load state of a processing unit may be inserted into a packet processed by said unit or a packet returned by a processing unit may be interpreted as a flag for a free resource.

Besides, should excess traffic occur it can be redirected to another load balancing unit, wherein said excess traffic would be defined upon the number of active processing units.

- 4 -

The method according to the present invention provides a more effective utilization of the media processing resources, since the resources are managed on the basis of an effective resource allocation of the whole network element instead of managing resource allocations of single processing units. Thus, the benefits of a statistical multiplexing can be exploited easily.

An additional benefit is the deterministic, i.e. optimal, processing time that a single packet always encounters, since a processing unit serves only one packet at a time.

Moreover, with the method according to the present invention, the processing delay of a received packet is always optimal and very constant. Consequently, a constant processing time minimizes unwanted jitter and other possible fluctuations of the traffic flow.

According to the present invention, the object is further solved by providing a device unit for serving and processing packets of a communication connection, comprising means adapted to inform a load state of said device to a balancing unit; and means adapted to obtain a state of said communication connection.

In this processing device unit, said obtaining means can be adapted to retrieve said communication connection state from a data storage or from a packet being under processing.

According to the present invention, the object is still further solved by providing a device unit for balancing a load of each of multiple processing units performing a packet switched communication connection, comprising: means for maintaining a load state of each of said

- 5 -

processing units; and means adapted to select a processing unit on the basis of a respective load state.

In this balancing device unit, a load state of a processing unit may be contained in a table. The state can be expressed as a Boolean state or as value which corresponds to the percentage of load.

Further, said selecting means can be adapted such that a processing unit is selected also on the basis of a parameter indicating the service profile supported by a respective processing unit. In this case, said parameter could be contained in a table.

As a modification, the load balancing device unit may further comprise means adapted to insert a communication connection state into a packet to be routed.

In a preferred embodiment, the processing units are comprised of multicore digital signal processing means having a shared data storage for all cores, whereby said device comprises a first level of load balancing for selecting a digital signal processing means and a second level of load balancing for selecting a single core.

As another modification, the load balancing device unit may further comprise means for redirecting excess traffic to another load balancing device unit according to the present invention, wherein said excess traffic is defined upon the number of active processing units.

Furthermore, a system adapted to perform the method according to the present invention and/or comprising one or more devices according to the present invention does also solve the object.

35

Brief Description of the Drawings

Further details and advantages of the present invention as well as further modifications thereof are apparent
5 from the detailed description of the preferred embodiments which are to be taken in conjunction with the appended drawings, in which:

Fig. 1 shows a first embodiment of the present invention;
10

Fig. 2 shows a second embodiment of the present invention; and

Fig. 3 shows a conventional arrangement.
15

Detailed Description of the preferred Embodiments

The present invention introduces a load balancing unit (or DSP selector) in front of the DSP resource pool. The
20 main idea behind the load balancing unit is to remove (or at least substantially reduce) the need for separate and poorly scaleable resource management layers.

The figures 1 and 2 present the two preferred
25 implementation options. It is common to both implementations that the processing unit is selected by the load balancing unit on a per-packet basis.

Specifically, in fig. 1, reference numeral 11 designates
30 1...N processing units for serving packets of a communication connection; reference numeral 12 designates the load balancing unit; reference numeral 14 designates a data storage; and reference numeral 13 designates a routing device for routing packets of a communication
35 connection.

- 7 -

According to the arrangement depicted in fig. 1, a connection state is stored in the data storage 14 by the load balancing unit 12 as will be apparent from the description given below.

Next, in fig. 2, reference numeral 21 designates 1...N processing units for serving packets of a communication connection; reference numeral 22 designates the load balancing unit; reference numeral 24 designates a shared data storage; and reference numeral 23 designates a routing device for routing packets of a communication connection.

According to the arrangement depicted in fig. 2, a connection state is stored in the shared data storage 24 by processing units 21 as will be apparent from the description given below.

What is common to both implementations is that the load balancing unit keeps track of the total utilization of the processing units and this overall load information can be provided for other network management processes. Specifically, the arrangement that is depicted in figures 1 and 2 provides a streamlined and more cost-effective way to manage the available resource pool. The main idea behind the present invention is that the single processing units are not dedicated to serve a specific connection (or a call). Instead, the load balancing unit selects any free processing unit on a per-packet basis. The current connection state is obtained from a data storage that may be located either at the load balancing unit (fig. 1) or at the processing units (as shared memory as is depicted in fig. 2). In the former case, it would be essential that the connection state is inserted

- 8 -

into the packets by the load balancing unit. The load balancing unit maintains the load state of each processing unit (preferably as a Boolean state) and selects any of the free (=non-active) processing units to
5 serve and process the received packet. The selection of a processing unit, e.g. in a round-robin fashion, results the automatic load balancing for the system. A supported service profile for each processing unit (e.g. only GSM codecs) may also be maintained and used as an additional
10 selection criteria.

Furthermore, the conveyance of the load state from each processing unit to the load balancing unit may happen either by a hardware based mechanism (such as dedicated
15 pin, shared memory etc.) or a packet based mechanism (such as inserting the load state to returning (processed) packets or just interpreting a returning packet as a flag for a free resource). The load balancing unit and the processing units may be interconnected for
20 example with Ethernet/IP, thus they do not require a physical co-location.

Some functionalities that the processing units, the load balancing unit and the data storage provide in preferred
25 embodiments of the invention are outlined in the following.

The processing unit implements a mechanism to inform the load status to the load balancing unit which can be a
30 hardware based mechanism (dedicated pin, shared memory etc.) or a packet based, e.g. inserting the status in the processed packets. Further, the processing unit comprises means for obtaining the connection state from the data storage or from the received packet.

35

- 9 -

The load balancing unit which can also be one of the processing units implements a table that contains the load status of each DSP unit in a Boolean format (free or not free) or as a percentage of load (0...100% load).

5 Further, it comprises means for selecting a resource based on the load status, wherein a parameter that indicates the supported service profile for each processing unit (e.g. only EFR codec) may also be used as an additional selection variable. Optionally, the load
10 balancing unit may also comprise means for inserting the connection state to a routed packet.

The data storage has to maintain the connection states by mapping them to suitable connection identifiers such as
15 UDP ports and may lock the states in order to handle bursts of packets.

While the above may be considered as a basic arrangement according to the present invention, further developments
20 of the same invention may be as follows.

Since the state-of-the art signal processors consist of multiple cores (e.g. 4-8) per one physical chip, one possibility could be to implement the load balancing unit
25 functionality inside each multi-core DSP device which usually have a shared memory (data storage) for all cores. This way, there would be two levels of load balancing: one for selecting the DSP and a second level for selecting a single core.

30

In addition, it would also be preferred to have a redirecting functionality. That is, the excess traffic of a first load balancing unit could be redirected to another load balancing unit if a certain limit (or load)

- 10 -

is exceeded, i.e. when all or most of the processing units are active when a new packet arrives.

According to the above, the benefits of a streamlined
5 resource management, an automated load balancing between
a high number of processing units, a more efficient
utilization of the DSP resources resulting in a
statistical multiplexing with a high number of processing
units managed as a whole, the guaranteeing of a
10 deterministic processing delay for each packet leading to
a minimum delay and a smooth traffic pattern, and
maintaining the possibility to still be able to dedicate
processing units to a specific service as an optimal
utilization of memory thus gaining a highest number of
15 channels can be achieved.

According to the above, it is provided a method of
balancing the load of resources in a packet switched
connection within a communication system, said system
20 comprising processing units 11; 21 for performing
communication, at least one load balancing unit 12; 22
for distributing the load to said processing units 11;
21, and a data storage 14; 24, said method comprising the
steps of: obtaining a current connection state as well as
25 a current load state of said processing units from said
data storage 14; 24; selecting by said load balancing
unit 12; 22 a processing unit on a per-packet basis; and
maintaining information about the load state of each
processing unit 11; 21 so that said selecting step is
30 performed by selecting a processing unit to serve and
process a respective packet based on the load state.

While it is described above what is presently considered
to be the preferred embodiments of the present invention,
35 it is apparent to those skilled in the art that various

- 11 -

modifications are possible to the present invention without departing from the spirit and scope thereof which is defined in the appended claims.

- 12 -

Claims

1. A method of balancing the load of resources in a packet switched connection within a communication system, said system comprising processing units (11; 21) for performing communication, at least one load balancing unit (12; 22) for distributing the load to said processing units (11; 21), and a data storage (14; 24), said method comprising the steps of:
- 10 obtaining a current connection state as well as a current load state of said processing units from said data storage (14; 24);
- selecting by said load balancing unit (12; 22) a processing unit on a per-packet basis; and
- 15 maintaining information about the load state of each processing unit (11; 21) so that said selecting step is performed by selecting a processing unit to serve and process a respective packet based on the load state.
- 20 2. A method according to claim 1, wherein said data storage is accessed to by said load balancing unit.
3. A method according to claim 1, wherein said data storage is accessed to by said processing units.
- 25 4. A method according to claim 1, wherein said information about the load state is maintained as a Boolean state.
- 30 5. A method according to claim 1, wherein a processing unit is selected in a round-robin fashion.
6. A method according to claim 1, wherein a supported service profile for each processing unit is maintained.

- 13 -

7. A method according to claim 6, wherein said supported service profile is used as additional selection criteria.

8. A method according to claim 1, wherein said load
5 balancing unit obtains a load state from each processing unit upon a hardware based mechanism.

9. A method according to claim 1, wherein said load
balancing unit obtains a load state from each processing
10 unit upon a packet based mechanism.

10. A method according to claim 9, wherein a load state of a processing unit is inserted into a packet processed by said unit.

15 11. A method according to claim 9, wherein a packet returned by a processing unit is interpreted as a flag for a free resource.

20 12. A method according to claim 1, wherein excess traffic is redirected to another load balancing unit, said excess traffic being defined upon the number of active processing units.

25 13. A device unit for serving and processing packets of a communication connection, comprising:

means adapted to inform a load state of said device to a balancing unit; and

means adapted to obtain a state of said
30 communication connection.

14. A device unit according to claim 13, wherein said obtaining means is adapted to retrieve said communication connection state from a data storage.

35

- 14 -

15. A device unit according to claim 13, wherein said obtaining means is adapted to retrieve said communication connection state from a packet being under processing.

- 5 16. A device unit for balancing a load of each of multiple processing units performing a packet switched communication connection, comprising:

means for maintaining a load state of each of said processing units; and

- 10 means adapted to select a processing unit on the basis of a respective load state.

17. A device according to claim 16, wherein a load state of a processing unit is contained in a table.

15

18. A device according to claim 16, wherein a load state of a processing unit is expressed as a Boolean value.

19. A device according to claim 16, wherein a load state
20 of a processing unit is expressed as value which corresponds to the percentage of load.

20. A device according to claim 16, wherein said selecting means is adapted such that a processing unit is
25 selected also on the basis of a parameter indicating the service profile supported by a respective processing unit.

21. A device according to claim 20, wherein said
30 parameter is contained in a table.

22. A device according to claim 16, further comprising means adapted to insert a communication connection state into a packet to be routed.

35

- 15 -

23. A device according to claim 16, wherein the processing units are comprised of multicore digital signal processing means having a shared data storage for all cores, whereby said device comprises a first level of load balancing for selecting a digital signal processing means and a second level of load balancing for selecting a single core.

24. A device according to claim 16, further comprising means for redirecting excess traffic to another device according to claim 16, wherein said excess traffic is defined upon the number of active processing units.

25. A system adapted to perform a method according to any of the claims 1 to 12.




FIG. 1




FIG. 2




FIG. 3
(PRIOR ART)

INTERNATIONAL SEARCH REPORT

Int onal Application No
PCT/IB 02/02492

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6 272 522 B1 (DOOLITTLE ROBERT ARTHUR ET AL) 7 August 2001 (2001-08-07) detail 34 abstract	1-25
X	EP 0 782 072 A (MITSUBISHI ELECTRIC CORP) 2 July 1997 (1997-07-02) column 4, line 9-37	1-25
X	US 2001/043585 A1 (HUMMEL HEINRICH) 22 November 2001 (2001-11-22) abstract; claims 1-13	1-25
X	US 5 655 120 A (HELD WALTER ET AL) 5 August 1997 (1997-08-05) column 6, line 35-40; claims 1-11	13, 15-19,24
	--- -/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the International filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

6 February 2003

Date of mailing of the international search report

20. 02. 2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

KRISTOFFER OGBJER/JA

INTERNATIONAL SEARCH REPORT

International Application No

PCT/IB 02/02492

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/064160 A1 (KIM JONG TAK) 30 May 2002 (2002-05-30) the whole document ---	13, 15-19,24
X	US 5 978 844 A (OURA TETSUO ET AL) 2 November 1999 (1999-11-02) claims 1-15 ---	13, 15-19,24
X	US 4 748 558 A (HIROSAWA TOSHIO ET AL) 31 May 1988 (1988-05-31) abstract; claims 1-12 ---	16-19, 23,24
A	WO 96 18149 A (ERICSSON TELEFON AB L M ;ANDERSSON SOFIA BIRGITTA (SE); HOLTE ROST) 13 June 1996 (1996-06-13) claim 4 ---	1-25
A	EP 1 022 658 A (SIEMENS AG) 26 July 2000 (2000-07-26) abstract; claims 1-26 ---	1-25
A	US 5 991 808 A (BRODER ANDREI ET AL) 23 November 1999 (1999-11-23) claims 1-36 ---	1-25
E	US 2002/087694 A1 (DAOUD RAJA ET AL) 4 July 2002 (2002-07-04) paragraph [0038] -----	1-25

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/IB 02/02492

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 6272522	B1	07-08-2001	AU 1725300 A WO 0030322 A2	05-06-2000 25-05-2000
EP 0782072	A	02-07-1997	JP 9179820 A DE 69619213 D1 DE 69619213 T2 EP 1074913 A2 EP 1054323 A2 EP 1054324 A2 EP 1152337 A1 EP 0782072 A1 US 6401121 B1	11-07-1997 21-03-2002 07-11-2002 07-02-2001 22-11-2000 22-11-2000 07-11-2001 02-07-1997 04-06-2002
US 2001043585	A1	22-11-2001	EP 1133112 A2	12-09-2001
US 5655120	A	05-08-1997	EP 0645702 A1 AT 195185 T CA 2132672 A1 DE 59310083 D1 DK 645702 T3 ES 2149794 T3 JP 7152591 A	29-03-1995 15-08-2000 25-03-1995 07-09-2000 06-11-2000 16-11-2000 16-06-1995
US 2002064160	A1	30-05-2002	BR 0105526 A CN 1356816 A	02-07-2002 03-07-2002
US 5978844	A	02-11-1999	JP 9135273 A	20-05-1997
US 4748558	A	31-05-1988	JP 61253572 A	11-11-1986
WO 9618149	A	13-06-1996	AT 202225 T AU 716774 B2 AU 4277196 A BR 9509895 A CA 2206373 A1 CN 1169191 A DE 69521337 D1 DE 69521337 T2 DK 796463 T3 EP 0796463 A2 ES 2161918 T3 FI 972407 A GR 3036045 T3 JP 10511785 T NO 972597 A WO 9618149 A2	15-06-2001 09-03-2000 26-06-1996 25-11-1997 13-06-1996 31-12-1997 19-07-2001 11-10-2001 03-09-2001 24-09-1997 16-12-2001 06-06-1997 28-09-2001 10-11-1998 05-08-1997 13-06-1996
EP 1022658	A	26-07-2000	EP 1022658 A1 WO 0043874 A1 EP 1149338 A1 JP 2002535771 T	26-07-2000 27-07-2000 31-10-2001 22-10-2002
US 5991808	A	23-11-1999	NONE	
US 2002087694	A1	04-07-2002	GB 2373673 A JP 2002245017 A	25-09-2002 30-08-2002